

Online Hyperparameter Optimization for Class-Incremental Learning

Yaoyao Liu¹ Yingying Li² Bernt Schiele¹ Qianru Sun³

¹Max Planck Institute for Informatics, Saarland Informatics Campus

²Computing and Mathematical Sciences, California Institute of Technology

³School of Computing and Information Systems, Singapore Management University

{yaoyao.liu, schiele}@mpi-inf.mpg.de yingli2@caltech.edu qianrusun@smu.edu.sg

Abstract

Class-incremental learning (CIL) aims to train a classification model while the number of classes increases phase-by-phase. An inherent challenge of CIL is the stability-plasticity tradeoff, i.e., CIL models should keep stable to retain old knowledge and keep plastic to absorb new knowledge. However, none of the existing CIL models can achieve the optimal tradeoff in different data-receiving settings—where typically the training-from-half (TFH) setting needs more stability, but the training-from-scratch (TFS) needs more plasticity. To this end, we design an online learning method that can adaptively optimize the tradeoff without knowing the setting as a priori. Specifically, we first introduce the key hyperparameters that influence the trade-off, e.g., knowledge distillation (KD) loss weights, learning rates, and classifier types. Then, we formulate the hyperparameter optimization process as an online Markov Decision Process (MDP) problem and propose a specific algorithm to solve it. We apply local estimated rewards and a classic bandit algorithm Exp3 [4] to address the issues when applying online MDP methods to the CIL protocol. Our method consistently improves top-performing CIL methods in both TFH and TFS settings, e.g., boosting the average accuracy of TFH and TFS by 2.2 percentage points on ImageNet-Full, compared to the state-of-the-art [23].¹

1. Introduction

Real-world problems are ever-changing, with new concepts and new data being continuously observed. Ideal AI systems should have the ability to learn new concepts from the new data, also known as *plasticity*, while maintaining the ability to recognize old concepts, also known as *stability*. However, there is a fundamental *tradeoff* between plasticity and stability: too much plasticity may re-

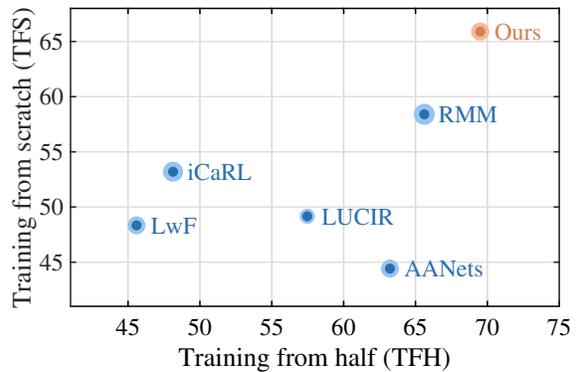


Figure 1. Average accuracy (%) on CIFAR-100 25-phase, using two data-receiving settings: 1) *training-from-half* (TFH): a large amount of data is available beforehand to pre-train the encoder; 2) *training-from-scratch* (TFS): classes come evenly in each phase. **Dark blue** and **orange** indicate the baselines and our method, respectively. Light-color circles are confidence intervals. Notice that methods with strong KD losses, e.g., LUCIR [14], AANets [22], and RMM [23], tend to provide worse performance in TFS than TFH, while methods with weak KD losses, e.g., iCaRL [30] and LwF [21], tend to provide worse performance in TFH than TFS. Our method uses an online learning algorithm to produce the key hyperparameters, e.g., the weights that control which KD losses are used. Thus, our method achieves the highest performance in both TFS and TFH.

sult in the *catastrophic forgetting* of old concepts, while too much stability restricts the ability to adapt to new concepts [25, 26, 29]. To encourage related research, Rebuffi et al. [30] defined the class-incremental learning (CIL) protocol, where the training samples of different classes come to the model phase-by-phase and most of the past data are removed from the memory.

Recently, many methods have been proposed to balance the stability-plasticity tradeoff for *different data-receiving settings* of CIL. For example, the strong feature knowl-

¹Code: <https://class-il.mpi-inf.mpg.de/online/>

edge distillation (KD) loss function is usually adopted when a large amount of data is available beforehand (e.g., the *training-from-half* (TFH) setting) since it encourages stability [14, 22, 23]; while the weak logit KD loss function is popular when data and classes are received evenly in each phase (e.g., the *training-from-scratch* (TFS) setting) since it provides more plasticity [7, 21, 30].

Most CIL algorithms pre-fix the tradeoff balancing methods, usually according to which data-receiving setting will be used in the experiments. However, in real-world scenarios, it is difficult to anticipate how data will be received in the future. Hence, a pre-fixed method is no longer proper for balancing the stability and plasticity of the actual data stream, thus generating worse performance. This can also be validated in Figure 1. Notice that the methods with weak KD, e.g., iCaRL [30] and LwF [21], provide worse performance in TFH than in TFS since weak KD provides too much plasticity for TFH, while the methods with strong KD, e.g., LUCIR [14], AANets [22], and RMM [23], perform worse in TFS than in TFH due to too much stability.

Therefore, a natural question is: how to design an *adaptive trade-off balancing method* to achieve good performance *without knowing* how data will be received *beforehand*? To tackle this, we propose an online-learning-inspired method to adaptively adjust key hyperparameters that affect the trade-off balancing performance in CIL. In our method, we introduce hyperparameters to control the choice of KD loss functions, learning rates, and classifier types, which are key algorithm choices that affect the tradeoff balancing performance.² In this way, deciding this choice is transformed into a hyperparameter optimization (HO) problem.

This HO problem cannot be directly solved because future data are not available. Thus, we borrow ideas from online learning, which is a widely adopted approach to adaptively tune the decisions without knowing the future data a priori while still achieving good performance in hindsight.

However, in CIL, our decisions affect not only the next phase but also all the future phases, which is different from the standard online learning setting [3]. To capture the dependence across phases, we formulate the HO problem in CIL as an online MDP, which is a generalized version of online learning. Further, we propose a new algorithm based on [10] to solve this online MDP problem. Our algorithm differs from the standard online MDP algorithm in [10] in two aspects:

- In CIL, we cannot directly observe the reward (i.e., validation accuracy) because the validation data is not ac-

²The impact of KD losses has been discussed before. Learning rates naturally affect how fast the model learns new concepts. We also adjust the classifier type because empirical results [14, 30] show that the nearest class mean (NCM) and fully-connected (FC) classifiers perform better under more plasticity and stability, respectively.

cessible during training. To address this issue, we estimate the reward by rebuilding local training and validation sets during policy learning in each phase, and computing the estimated reward on the local validation sets.

- In CIL, we only have access to the model generated by the selected hyperparameters instead of the models generated by other hyperparameters. In other words, we only have bandit feedback instead of full feedback as assumed in [10]. To address this, we revise the algorithm in [10] by combining it with a classic bandit algorithm, Exp3 [4].

Empirically, we find our method performs well consistently. We conduct extensive CIL experiments by plugging our method into three top-performing methods (LUCIR [14], AANets [22], and RMM [23]) and testing them on three benchmarks (i.e., CIFAR-100, ImageNet-Subset, and ImageNet-Full). Our results show the consistent improvements of the proposed method, e.g., boosting the average accuracy of TFH and TFS by 2.2 percentage points on ImageNet-Full, compared to the state-of-the-art [23].

Lastly, it is worth mentioning that our method can also be applied to optimize other key hyperparameters in CIL, e.g., memory allocation [23].

Summary of our contributions. Our contributions are three-fold: 1) an online MDP formulation that allows online updates of hyperparameters that affect the balance of the stability and plasticity in CIL; 2) an Exp3-based online MDP algorithm to generate adaptive hyperparameters using bandit and estimated feedback; 3) extensive comparisons and visualizations for our method in three CIL benchmarks, taking top-performing methods as baselines.

2. Related Work

Class-incremental learning (CIL). There are three main lines of work to address the stability-plasticity trade-off in CIL. *Distillation-based* methods [8, 15, 16, 21, 27, 37, 41] introduce different knowledge distillation (KD) losses to consolidate previous knowledge when training the model on new data. The key idea is to enforce model prediction logits [21, 30], feature maps [8, 14], or topologies in the feature space [36] to be close to those of the pre-phase model. *Memory-based* methods [7, 24, 28, 30, 34] preserve a small number of old class data (called exemplars) and train the model on them together with new class data. *Network-architecture-based* methods [1, 32, 38–40] design incremental network architectures by expanding the network capacity for new data or freezing partial network parameters to keep the knowledge of old classes. However, all the above methods are either too stable or too plastic to perform well in both TFS and TFH settings. Our method learns an online policy to generate hyperparameters that balance stabil-

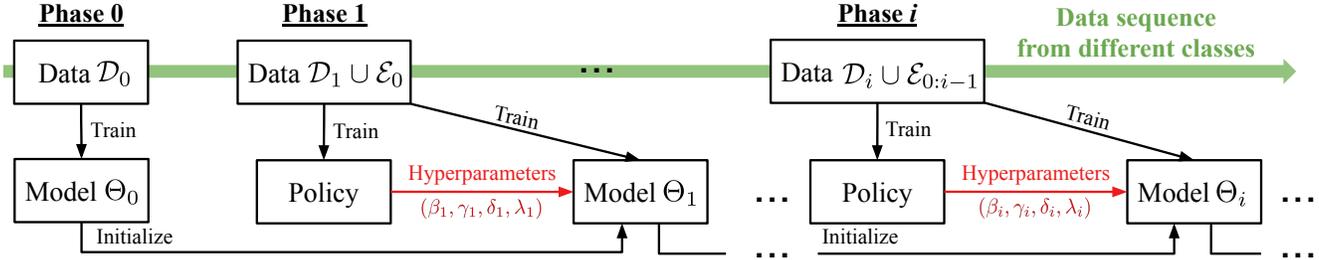


Figure 2. The computing flow of our method. We formulate the CIL task as an online MDP: each phase in CIL is a stage in the MDP, and the CIL models are the states. We train the policy to produce actions, which contain the hyperparameters we use in the CIL training. We illustrate the training process of each phase in Figure 3.

ity and plasticity. Therefore, our method performs well in both TFS and TFH settings.

Reinforcement learning (RL) aims to learn a policy in an environment, which is typically formulated as an MDP. Some CIL papers also deploy RL algorithms in their frameworks. Xu et al. [39] used RL to expand its backbone network when a new task arrives adaptively. Liu et al. [23] used RL to learn a policy to adjust the memory allocation between old and new class data dynamically along with the learning phases. Our method focuses on learning a policy to produce key hyperparameters. Besides, the existing methods need to solve the complete MDP, which is time-consuming. Here, we formulate the CIL task as an online MDP. Thus, our method is more time-efficient.

Online learning observes a stream of samples and makes a prediction for each element in the stream. There are mainly two settings in online learning: full feedback and bandit feedback. *Full feedback* means that the full reward function is given at each stage. It can be solved by Best-Expert algorithms [9]. *Bandit feedback* means that only the reward of the implemented decision is revealed. If the rewards are independently drawn from a fixed and unknown distribution, we may use e.g., Thompson sampling [2] and UCB [5] to solve it. If the rewards are generated in a non-stochastic version, we can solve it by e.g., Exp3 [4]. *Online MDP* is an extension of online learning. Many studies [10, 18–20] aim to solve it by converting it to online learning. In our case, we formulate the CIL as an online MDP and convert it into a classic online learning problem. The rewards in our MDP are non-stochastic because the training and validation data change in each phase. Therefore, we design our algorithm based on Exp3 [4].

Hyperparameter optimization (HO). There are mainly two popular lines of HO methods: gradient-based and meta-learning-based. Gradient-based HO methods [6] make it possible to tune the entire weight vectors associated with a neural network layer as hyperparameters. Meta-learning-based HO methods [11] use a bilevel program to optimize the hyperparameters. However, all these methods only consider time-invariant environments. Our online method

learns hyperparameters that adapt to the time-varying environments in CIL.

3. Preliminaries

Class-incremental learning (CIL). The general CIL pipeline is as follows. There are multiple phases during which the number of classes gradually increases to the maximum [8, 14, 15, 24]. In the 0-th phase, we observe data \mathcal{D}_0 , and use it to learn an initial model Θ_0 . After this phase, we can only store a small subset of \mathcal{D}_0 (i.e., exemplars denoted as \mathcal{E}_0) in memory used as replay samples in later phases. In the i -th phase ($i \geq 1$), we get new class data \mathcal{D}_i and load exemplars $\mathcal{E}_{0:i-1} = \mathcal{E}_0 \cup \dots \cup \mathcal{E}_{i-1}$ from the memory. Then, we initialize Θ_i with Θ_{i-1} , and train it using $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$. We evaluate the model Θ_i on a test set $\mathcal{Q}_{0:i}$ for all classes observed so far. After that, we select exemplars $\mathcal{E}_{0:i}$ from $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$ and save them in the memory.

Existing work mainly focus on two data-receiving settings: *training-from-half* (TFH) [14] and *training-from-scratch* (TFS) [30] settings. In TFH, we are given half of all classes in the 0-th phase, and observe the remaining classes evenly in the subsequent N phases. In TFS, we are given the same number of classes in all N phases.

4. Methodology

As illustrated in Figures 2 and 3, we formulate CIL as an online MDP and learn a policy to produce the hyperparameters in each phase. In this section, we introduce the online MDP formulation, show optimizable hyperparameters, and provide an online learning algorithm to train the policy. Algorithm 1 summarizes the overall training steps of the proposed method. Algorithm 2 presents the training and evaluation process for a specific action.

4.1. An Online MDP Formulation for CIL

Optimizing hyperparameters in CIL should be online inherently: training and validation data changes in each phase, so the hyperparameters should be adjusted accordingly. Thus, it is intuitive to formulate the CIL as an online

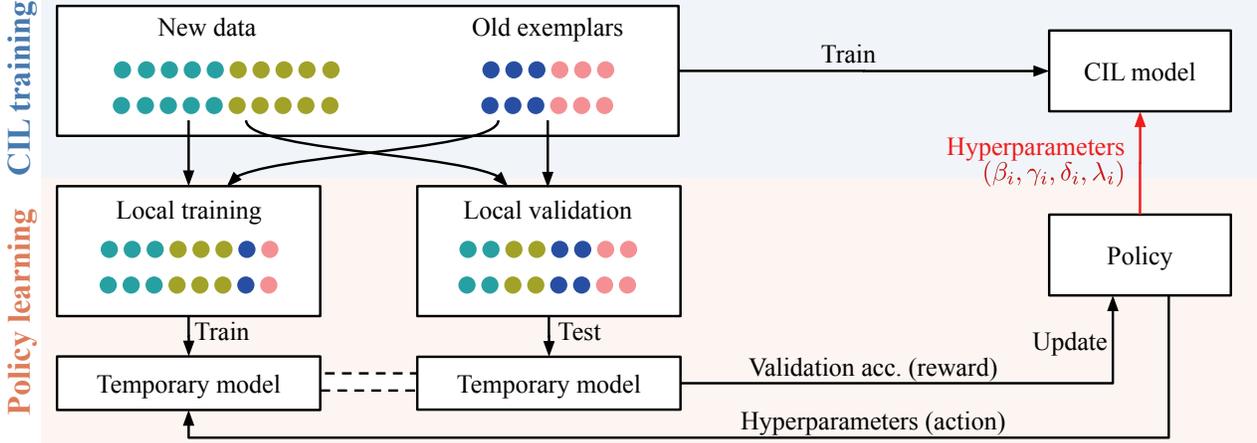


Figure 3. The training process of our online learning method in the i -th phase. It includes *policy learning* and *CIL training*. **(a) Policy learning.** 1) We construct a class-balanced subset from all training data as the local validation set and use the remaining data as the local training set. 2) We initialize the temporary model with Θ_{i-1} . 3) We sample an action using the current policy and deploy the hyperparameters on the temporary model according to the action. 4) We train it on the local training set for M_1 epochs, and evaluate it on the local test set. 5) We use the validation accuracy as the reward and update the policy. We update the policy for T iterations by repeating Steps 2-5. **(b) CIL training.** We sample an action using the learned policy and deploy the hyperparameters on the CIL model. Then, we train the CIL model on all training data for M_2 epochs. We set $M_1 = \lceil 0.1M_2 \rceil$ to speed up the policy learning.

MDP. In the following, we provide detailed formulations.

Stages. Each phase in the CIL task can be viewed as a stage in the online MDP.

States. The state should define the current situation of the intelligent agent. In CIL, we use the model Θ_i as the state of the i -th phase/stage. We use \mathbb{S} to denote the state space.

Actions. We use a vector consisting of the hyperparameters in the i -th phase as the action \mathbf{a}_i . When we take the action \mathbf{a}_i , we deploy the corresponding hyperparameters. We denote the action space as \mathbb{A} . Please refer to the next subsection, *Optimizable Hyperparameters*, for more details.

Policy $\mathbf{p} = \{p(\mathbf{a}|\Theta_i)\}_{\mathbf{a} \in \mathbb{A}}$ is a probability distribution over the action space \mathbb{A} , given the current state Θ_i .

Environments. We define the training and validation data in each phase as the environment. In the i -th phase, the environment is $\mathcal{H}_i = (\mathcal{E}_{0:i-1} \cup \mathcal{D}_i, \mathcal{Q}_{0:i})$, where $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$ is the training data and $\mathcal{Q}_{0:i}$ is the corresponding validation data. The environment is time-varying because we are given different training and validation data in each phase.

Rewards. CIL aims to train a model that is efficient in recognizing all classes seen so far. Therefore, we use the validation accuracy as the reward in each phase. Our objective is to maximize a cumulative reward, i.e., $R = \sum_{i=1}^N r_{\mathcal{H}_i}(\Theta_i, \mathbf{a}_i)$, where $r_{\mathcal{H}_i}(\Theta_i, \mathbf{a}_i)$ denotes the i -th phase reward, i.e., the validation accuracy of Θ_i . The reward function $r_{\mathcal{H}_i}$ changes with \mathcal{H}_i , so it is time-varying.

4.2. Optimizable Hyperparameters

In this part, we introduce the optimizable hyperparameters, and how to define the actions and action space based on

the hyperparameters. We consider three kinds of hyperparameters that significantly affect stability and plasticity: 1) KD loss weights, 2) learning rates, and 3) classified types.

1) KD loss weights. We first introduce two KD losses (i.e., logit and feature KD losses) and then show how to use the KD loss weights to balance them.

Logit KD loss is proposed in [13] and widely applied in CIL methods [21, 22, 30]. Its motivation is to make the current model Θ_i mimic the prediction logits of the old model Θ_{i-1} :

$$\mathcal{L}_{\text{logit}} = - \sum_{k=1}^K \eta_k(\mu(x; \Theta_{i-1})) \log \eta_k(\mu(x; \Theta_i)), \quad (1)$$

where $\mu(x; \Theta)$ is a function that maps the input mini-batch (x, y) to the prediction logits using the model Θ . $\eta_k(v) = v_k^{1/\tau} / \sum_j v_j^{1/\tau}$ is a re-scaling function for the k -th class prediction logit and τ is a scalar set to be greater than 1.

Feature KD loss [8, 14] aims to enforce a stronger constraint on the previous knowledge by minimizing the cosine similarity between the features from the current model Θ_i and the old model Θ_{i-1} . It can be computed as follows,

$$\mathcal{L}_{\text{feat}} = 1 - S_c(f(x; \Theta_i), f(x; \Theta_{i-1})), \quad (2)$$

where $f(x; \Theta)$ denotes a function that maps the input image x to the features using the model Θ . $S_c(v_1, v_2)$ denotes the cosine similarity between v_1 and v_2 .

The overall loss in the i -th phase is a weighted sum of the classification and different KD losses:

$$\mathcal{L}_{\text{overall}} = \mathcal{L}_{\text{CE}} + \beta_i \mathcal{L}_{\text{logit}} + \gamma_i \mathcal{L}_{\text{feat}}, \quad (3)$$

where β_i and γ_i are the weights of the logit and feature KD losses, respectively. \mathcal{L}_{CE} is the standard cross-entropy classification loss. Existing methods [14, 21–23, 30] can be viewed as using fixed heuristic KD loss weights, e.g., $\beta_i \equiv 1$ and $\gamma_i \equiv 0$ in iCaRL [21, 30]. Instead, our method optimizes β_i and γ_i online. Thus, we can balance the model’s stability and plasticity by adjusting β_i and γ_i . We apply different weights for the logit and feature KD losses so that we can achieve fine-grained control over the intensity of knowledge distillation.

2) Learning rate is another important hyperparameter that affects the model’s stability and plasticity. We empirically find that a lower learning rate makes the CIL model more stable, while a higher learning rate makes the CIL model more plastic. If we use λ_i to denote the learnable learning rate in the i -th phase, we update the CIL model as follows,

$$\Theta_i \leftarrow \Theta_i - \lambda_i \nabla_{\Theta_i} \mathcal{L}_{\text{overall}}. \quad (4)$$

Another hyperparameter, the number of training epochs, has similar properties to the learning rate. We choose to optimize the learning rate and fix the number of epochs because it empirically works better with our online learning algorithm.

3) Classifier type. Motivated by the empirical analysis, we consider two classifier types in our study: *nearest class mean* (NCM) [30, 35] and *fully-connected* (FC) [14, 24] classifiers. For the NCM classifier, we first compute the mean feature for each class using the new data and old exemplars. Then we perform a nearest neighbor search using the Euclidean distance on the L_2 normalized mean features to get the final predictions. It is observed empirically that the NCM classifier tends to work better on the models with high plasticity, while the FC classifier performs better on the models with high stability [14, 30]. Thus, we propose to use a hyperparameter, classifier type indicator δ_i , to control the final predictions during the evaluation:

$$\mu(x; \Theta_i) = \mu_{\text{ncm}}(x; \Theta_i)[\delta_i = 1] + \mu_{\text{fc}}(x; \Theta_i)[\delta_i = 0], \quad (5)$$

where $\delta_i \in \{0, 1\}$, μ_{ncm} and μ_{fc} are the predictions on the input image x using the NCM and FC classifiers, respectively.

Summary: actions and action space. In summary, we define the action as $\mathbf{a}_i = (\beta_i, \gamma_i, \lambda_i, \delta_i)$, which consists of the following hyperparameters: KD loss weights β_i and γ_i , learning rate λ_i , and classifier type indicator δ_i . For the hyperparameters that may vary in a continuous range, we *discretize* them to define a *finite* action space.³ In the next subsection, we show how to learn the policy in each phase.

³Though discretization suffers the curse of dimensionality, our experiments show that with a coarse grid, we already have significant improvements over pre-fixed hyperparameters.

4.3. Policy Learning

A common approach to solving an online MDP is to approximate it as an online learning problem and solve it using online learning algorithms [2, 4, 9]. We also take this approach, and our approximation follows [10], which achieves the optimal regret. In their paper, Even-Dar et al. [10] relax the Markovian assumption of the MDP by decoupling the cumulative reward function and letting it be time-dependent so that they can solve the online MDP by standard online learning algorithms. Such a decoupling requires the following assumptions. 1) Fast mixing: in CIL, the hyperparameters in an early phase do not have much impact on the test accuracy of the classes observed in the current phase. 2) The algorithm changes the hyperparameters slowly (this can be observed in *Experiments* & Figure 5). Thus, these assumptions fit our CIL problem.

However, we cannot directly apply the algorithms proposed in [10] to our problem. It is because their paper assumes *full feedback*, i.e., we can observe the rewards of all actions in each phase. Therefore, its online learning problem could be solved by Best Expert algorithms [9]. In CIL, we cannot observe any reward (i.e., validation accuracy) because the validation data is not accessible during training. To address this issue, we rebuild the local training and validation sets in each phase. In this way, our problem has *bandit feedback*: we can compute the reward of the implemented action. Therefore, we can solve our online learning problem based on Exp3 [4], a famous bandit algorithm.

In the following, we show how to rebuild the local training and validation sets, compute the decoupled cumulative reward, and learn the policy with Exp3.

Rebuilding local datasets. In the i -th phase, we need to access the validation set $\mathcal{Q}_{0:i}$ to compute the reward (i.e., the validation accuracy). However, we are not allowed to use $\mathcal{Q}_{0:i}$ during training because it violates the CIL benchmark protocol. Therefore, we replace $\mathcal{Q}_{0:i}$ with a class-balanced subset $\mathcal{B}_{0:i}$ sampled from the training data $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$. $\mathcal{B}_{0:i}$ contains the same number of samples for both the old and new classes. In this way, we can rebuild the local training and validation sets, and obtain the local environment $h_i = ((\mathcal{E}_{0:i-1} \cup \mathcal{D}_i) \setminus \mathcal{B}_{0:i}, \mathcal{B}_{0:i})$.

Decoupled cumulative reward. We create the decoupled cumulative reward function \hat{R} based on the original cumulative reward function $R = \sum_{j=1}^N r_{\mathcal{H}_j}(\Theta_j, \mathbf{a}_j)$. In the i -th phase, we compute \hat{R} as follows,

$$\hat{R}(\mathbf{a}_i, h_i) = \underbrace{\sum_{j=1}^{i-1} r_{\mathcal{H}_j}(\Theta_j, \mathbf{a}_j)}_{\text{Part I}} + \underbrace{\sum_{j=i}^{i+n} r_{h_i}(\Theta_j, \mathbf{a}_j)}_{\text{Part II}}, \quad (6)$$

where Part I is the historical rewards from the 1-st phase to the $(i-1)$ -th phase. It is a constant and doesn’t influence

policy optimization. Part II is the long-term reward of a time-invariant local MDP based on the local environment h_i . We use Part II as an estimation of the future rewards, following [10]. Because we don't know the total number of phases N during training, we assume there are n phases in the future. Furthermore, we fix the action \mathbf{a}_i in Part II to simplify the training process. Thus, \hat{R} can be reviewed as a function of \mathbf{a}_i and h_i .

Training policy with Exp3. Exp3 [4] introduces an auxiliary variable $\mathbf{w} = \{w(\mathbf{a})\}_{\mathbf{a} \in \mathbb{A}}$. After updating \mathbf{w} , we can determine the policy $\mathbf{p} = \{p(\mathbf{a}|\Theta_i)\}_{\mathbf{a} \in \mathbb{A}}$ by

$$\mathbf{p} = \mathbf{w} / \|\mathbf{w}\|. \quad (7)$$

The updating rule of \mathbf{w} is provided below.

In the 1-st phase, we initialize \mathbf{w} as $\{1, \dots, 1\}$. In each phase i ($i \geq 1$), we update \mathbf{w} for T iterations. In the t -th iteration, we sample an action $\mathbf{a}_t \sim \mathbf{p}$, apply the action \mathbf{a}_t to the CIL system, and compute the decoupled cumulative reward $\hat{R}(\mathbf{a}_t, h_i)$ using Eq. 6. After that, we update $w(\mathbf{a}_t)$ in \mathbf{w} as,

$$w(\mathbf{a}_t) \leftarrow w(\mathbf{a}_t) \exp(\xi \hat{R}(\mathbf{a}_t, h_i) / p(\mathbf{a}_t | \Theta_i)), \quad (8)$$

where ξ is a constant, which can be regarded as the learning rate in Exp3.

5. Experiments

We evaluate the proposed method on three CIL benchmarks, incorporate our method into three top-performing baseline methods, and boost their performances consistently in all settings. Below we describe the datasets and implementation details, followed by the results and analyses.

5.1. Datasets and Implementation Details

Datasets. We employ CIFAR-100 [17], ImageNet-Subset [30] (100 classes), and ImageNet-Full [31] (1000 classes) as the benchmarks. We use the same data splits and class orders as the related work [22, 23, 30].

Network architectures. We use a modified 32-layer ResNet [30] for CIFAR-100 and an 18-layer ResNet [12] for ImageNet, following [14, 22, 30]. We deploy the AANets [22] for the experiments based on AANets and RMM [23]. Further, we use a cosine normalized classifier without bias terms as the FC classifier, following [14, 24].

Configurations. We discretize the hyperparameter search space into 50 actions, i.e., $\text{card}(\mathbb{A})=50$. We update the policy for 25 iterations in each phase, i.e., $T=25$. For other configurations, we follow the corresponding baselines.

5.2. Results and Analyses

Tables 1 and 2 present the results of top-performing baselines w/ and w/o our method and some recent related

Algorithm 1: Our CIL algorithm in Phase i ($i \geq 1$)

Input : Old model Θ_{i-1} , training data $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$, validation data $\mathcal{Q}_{0:i}$, learnable parameters \mathbf{w} , numbers of epochs M_1 and M_2 .
Output: New model Θ_i , new exemplars $\mathcal{E}_{0:i}$, learnable parameters \mathbf{w} .

```

1 // Policy learning
2 if  $i=1$  then
3   Initialize  $\mathbf{w} = \{1, \dots, 1\}$ ;
4 for  $t$  in  $i, \dots, T$  do
5   Randomly sample a class-balanced subset  $\mathcal{B}_{0:i}$ 
   from  $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$ ;
6   Create the local environment
    $h_i = ((\mathcal{E}_{0:i-1} \cup \mathcal{D}_i) \setminus \mathcal{B}_{0:i}, \mathcal{B}_{0:i})$ ;
7   Set the policy  $\mathbf{p} = \mathbf{w} / \|\mathbf{w}\|$ ;
8   Sample an action  $\mathbf{a}_t \sim \mathbf{p}$ ;
9   for  $j$  in  $i, \dots, i+n$  do
10    Train  $\Theta_j$  for  $M_1$  epochs by Algorithm 2
    with inputs  $\Theta_{j-1}, \mathbf{a}_t, h_i$ ;
11    Collect the reward  $r_{h_i}(\Theta_j, \mathbf{a}_t)$ ;
12    Compute the cumulative reward  $\hat{R}(\mathbf{a}_t, h_i)$  by
    Eq. 6;
13    Update  $\mathbf{w}$  by Eq. 8;
14 // CIL training
15 Sample an action  $\mathbf{a}_i \sim \mathbf{p}$ ;
16 Train  $\Theta_i$  for  $M_2$  epochs by Algorithm 2 with inputs
    $\Theta_{i-1}, \mathbf{a}_i, \mathcal{H}_i = (\mathcal{E}_{0:i-1} \cup \mathcal{D}_i, \mathcal{Q}_{0:i})$ ;
17 Select new exemplars  $\mathcal{E}_{0:i}$  from  $\mathcal{E}_{0:i-1} \cup \mathcal{D}_i$  using
   herding [30].
```

Algorithm 2: Training and evaluation for action \mathbf{a}

Input : Old model Θ_{old} , action $\mathbf{a} = \{\beta, \gamma, \delta, \lambda\}$, environment $h = \{\mathcal{T}, \mathcal{Q}\}$.
Output: New model Θ , reward $r_h(\Theta, \mathbf{a})$ (i.e., the validation accuracy).

```

1 Initialize  $\Theta$  with  $\Theta_{\text{old}}$ ;
2 for epochs do
3   for mini-batch  $(x, y) \in \mathcal{T}$  do
4     Compute the loss  $\mathcal{L}_{\text{overall}}(x, y; \Theta_{\text{old}}, \beta, \gamma)$  by
     Eq. 3;
5     Update  $\Theta$  by Eq. 4 with  $\lambda$ ;
6   for mini-batch  $(x_{\text{val}}, y_{\text{val}}) \in \mathcal{Q}$  do
7     Compute predictions  $\mu(x_{\text{val}}; \Theta)$  by Eq. 5 with  $\delta$ ;
8     Compute the reward  $r_h(\Theta, \mathbf{a})$  using
      $\{(\mu(x_{\text{val}}), y_{\text{val}}))\}_{(x_{\text{val}}, y_{\text{val}}) \in \mathcal{Q}}$ .
```

work. Table 3 summarizes the results in seven ablative settings. Figure 4 compares the activation maps (using Grad-CAM [33]) produced by different methods in TFH and TFS. Figure 5 shows the values of hyperparameters produced by

| Methods | CIFAR-100, $N=5$ | | | CIFAR-100, $N=25$ | | | ImageNet-Subset, $N=5$ | | | ImageNet-Subset, $N=25$ | | |
|-------------|-----------------------|-----------------------|-----------------------|-----------------------|----------------|-----------------------|------------------------|-----------------------|-----------------------|-------------------------|-----------------------|-----------------------|
| | TFH | TFS | Avg. | TFH | TFS | Avg. | TFH | TFS | Avg. | TFH | TFS | Avg. |
| iCaRL [30] | 58.1 | 64.0 | 61.0 | 48.1 | 53.2 | 50.7 | 65.3 | 70.4 | 67.9 | 53.0 | 53.5 | 53.3 |
| PODNet [8] | 64.7 | 63.6 | 64.2 | 60.3 | 45.3 | 52.8 | 64.3 | 58.9 | 61.6 | 68.3 | 39.1 | 53.7 |
| DER [40] | 67.6 | 72.3 | 70.0 | 65.5 | 67.3 | 66.4 | 78.4 | 76.9 | 77.7 | 75.4 | 71.0 | 73.2 |
| FOSTER [38] | 70.4 | 72.5 | 71.5 | 63.8 | 70.7 | 67.3 | 80.2 | 78.3 | 79.3 | 69.3 | 72.9 | 71.1 |
| LUCIR [14] | 63.1 \pm 0.7 | 63.0 \pm 0.6 | 63.1 \pm 0.7 | 57.5 \pm 0.4 | 49.2 \pm 0.5 | 53.4 \pm 0.5 | 65.3 \pm 0.6 | 66.7 \pm 0.5 | 66.0 \pm 0.6 | 61.4 \pm 0.7 | 46.2 \pm 0.8 | 53.8 \pm 0.8 |
| w/ ours | 63.9 \pm 0.6 | 64.9 \pm 0.5 | 64.4 \pm 0.6 | 59.3 \pm 0.5 | 52.4 \pm 0.5 | 55.9 \pm 0.5 | 70.6 \pm 0.7 | 68.4 \pm 0.6 | 69.5 \pm 0.7 | 62.9 \pm 0.6 | 54.1 \pm 0.6 | 58.5 \pm 0.6 |
| | \uparrow 0.8 | \uparrow 1.9 | \uparrow 1.3 | \uparrow 1.8 | \uparrow 3.2 | \uparrow 2.5 | \uparrow 5.3 | \uparrow 1.7 | \uparrow 3.5 | \uparrow 1.5 | \uparrow 7.9 | \uparrow 4.7 |
| AANets [22] | 65.3 \pm 0.4 | 63.1 \pm 0.3 | 64.2 \pm 0.4 | 63.2 \pm 0.3 | 44.4 \pm 0.4 | 53.8 \pm 0.4 | 77.0 \pm 0.7 | 68.9 \pm 0.6 | 73.0 \pm 0.7 | 72.2 \pm 0.6 | 60.7 \pm 0.5 | 66.5 \pm 0.6 |
| w/ ours | 67.0 \pm 0.3 | 65.1 \pm 0.3 | 66.1 \pm 0.3 | 64.1 \pm 0.4 | 50.3 \pm 0.5 | 57.2 \pm 0.5 | 77.3 \pm 0.6 | 70.6 \pm 0.5 | 74.0 \pm 0.6 | 72.9 \pm 0.5 | 64.8 \pm 0.5 | 68.9 \pm 0.5 |
| | \uparrow 1.7 | \uparrow 2.0 | \uparrow 1.9 | \uparrow 0.9 | \uparrow 5.9 | \uparrow 3.4 | \uparrow 0.3 | \uparrow 1.7 | \uparrow 1.0 | \uparrow 0.7 | \uparrow 4.1 | \uparrow 2.4 |
| RMM [23] | 67.6 \pm 0.7 | 70.4 \pm 0.8 | 69.0 \pm 0.8 | 65.6 \pm 0.6 | 58.4 \pm 0.6 | 62.0 \pm 0.6 | 79.5 \pm 0.2 | 80.5 \pm 0.3 | 80.0 \pm 0.3 | 75.0 \pm 0.3 | 71.6 \pm 0.3 | 73.3 \pm 0.3 |
| w/ ours | 70.8 \pm 0.7 | 72.7 \pm 0.6 | 71.8 \pm 0.7 | 69.5 \pm 0.8 | 65.9 \pm 0.7 | 67.7 \pm 0.8 | 81.0 \pm 0.3 | 82.2 \pm 0.4 | 81.6 \pm 0.4 | 76.1 \pm 0.2 | 73.2 \pm 0.4 | 74.7 \pm 0.3 |
| | \uparrow 3.2 | \uparrow 2.3 | \uparrow 2.8 | \uparrow 3.9 | \uparrow 7.5 | \uparrow 5.7 | \uparrow 1.5 | \uparrow 1.7 | \uparrow 1.6 | \uparrow 1.1 | \uparrow 1.6 | \uparrow 1.4 |

Table 1. Average accuracy (%) across all phases on CIFAR-100 and ImageNet-Subset. The first block shows some recent CIL methods. The second block shows three top-performing baselines [14, 22, 23] w/ and w/o our method plugged in. “TFH” and “TFS” denote the *training-from-half* and *training-from-scratch* settings, respectively. “Avg.” shows the average of the “TFH” and “TFS” results. For “AANets” [22], we use its version based on PODNet [8]. We rerun the baselines using their open-source code in a unified setting for a fair comparison.

| Methods | ImageNet-Full, $N=5$ | | |
|------------|--------------------------------------|--------------------------------------|--------------------------------------|
| | TFH | TFS | Avg. |
| LUCIR [14] | 64.5 \pm 0.3 | 62.7 \pm 0.4 | 62.0 \pm 0.4 |
| w/ ours | 65.8 \pm 0.3 \uparrow 1.3 | 66.1 \pm 0.3 \uparrow 3.4 | 66.0 \pm 0.3 \uparrow 2.4 |
| RMM [23] | 69.0 \pm 0.5 | 66.1 \pm 0.4 | 67.6 \pm 0.5 |
| w/ ours | 70.7 \pm 0.5 \uparrow 1.7 | 68.9 \pm 0.5 \uparrow 2.8 | 69.8 \pm 0.5 \uparrow 2.2 |

Table 2. Average accuracy (%) on ImageNet-Full.

our method.

Comparison with the state-of-the-art. Tables 1 and 2 show that taking our method as a plug-in module for the state-of-the-art [23] and other baselines [14, 22] consistently improves their performance. For example, RMM [23] w/ ours gains 4.3 and 2.2 percentage points on CIFAR-100 and ImageNet-Full, respectively. Interestingly, we find that we can surpass the baselines more when the number of phases N is larger. E.g., on CIFAR-100, our method improves RMM by 5.7 percentage points when $N=25$, while this number is 2.8 percentage points when $N=5$. Our explanation is that the forgetting problem is more serious when the number of phases is larger. Thus, we need better hyperparameters to balance stability and plasticity.

Ablation study. Table 3 concerns eight ablation settings, and shows the results in both TFH and TFS for different numbers of phases ($N=5/25$). The detailed analyses are as follows.

1) **First block.** Row 1 shows the baseline [14].

2) **Second block: optimizable hyperparameters.** In our study, we optimize three kinds of hyperparameters that affect the model’s stability and plasticity: KD loss weights (β, γ) , learning rate λ , and classifier type indicator δ . Comparing Row 2 to Row 1, we can observe that optimizing

| No. | Optimizing | | | $N=5$ | | $N=25$ | |
|-----|-------------------|----------|-----------|--------------|--------------|--------------|--------------|
| | (β, γ) | δ | λ | TFH | TFS | TFH | TFS |
| 1 | Baseline | | | 63.11 | 62.96 | 57.47 | 49.16 |
| 2 | ✓ | | | 63.20 | 63.60 | 58.27 | 50.91 |
| 3 | ✓ | ✓ | | 63.23 | 64.08 | 58.20 | 51.94 |
| 4 | ✓ | ✓ | ✓ | 63.88 | 64.92 | 59.27 | 52.44 |
| 5 | Cross-val fixed | | | 63.33 | 64.02 | 57.50 | 51.64 |
| 6 | Offline RL [23] | | | 63.42 | 63.88 | 58.12 | 51.53 |
| 7 | Bilevel HO [11] | | | 63.20 | 63.02 | 57.56 | 49.42 |

Table 3. Ablation results (average accuracy %) on CIFAR-100. (β, γ) are KD loss weights. λ and δ denote learning rates and classifier types, respectively. The baseline is LUCIR [14]. Row 4 shows our best result.

the KD loss weights boosts the TFS accuracy more significantly. It is because the baseline, LUCIR [14], applies a strong regularization term (i.e., feature KD loss) which harms the TFH performance. Our method changes the regularization term by adjusting the KD loss weights, so it achieves better performance. Comparing Row 3 to Row 2, we can see that optimizing the classifier type indicator δ performs further improvements, especially in TFS. It is because the baseline deploys an FC classifier by default while our method learns to switch between different classifiers in different settings. Comparing Row 4 to Row 3, we can see the effectiveness of optimizing the learning rates in CIL. In summary, it is impressive that optimizing three kinds of hyperparameters together achieves the best results.

3) **Third block: hyperparameter learning methods.** For Row 5, we use cross-validation (i.e., all past, future, and validation data are accessible) to find a set of fixed hyperparameters and apply them to all phases. We can see that Row

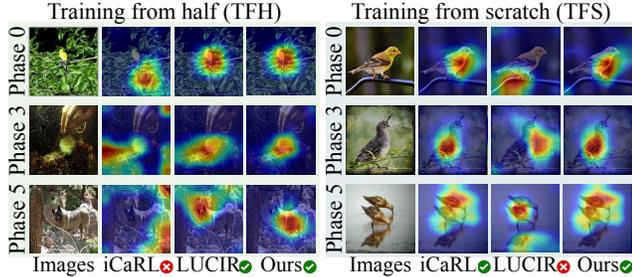


Figure 4. The activation maps using Grad-CAM [33] for the last phase model on ImageNet-Subset 5-phase. Samples are selected from the classes coming in the 0-th, 3-rd, and 5-th phases. Green ticks mean successful activation of discriminative features on object regions, while red crosses mean unsuccessful.

5 results are consistently lower than ours in Row 4, although it can access more data. It shows that we need to update the hyperparameters online in different phases. For Row 6, we use the policy pre-trained in the 0-th phase of the target CIL task using the framework proposed in [23] (compared to ours, it is offline). Comparing Row 6 with Row 4, we are happy to see that our online learning algorithm achieves better performance than the offline RL while we use much less training time (see the analysis in the supplementary). For Row 7, we use the bilevel hyperparameter optimization method [11]. Comparing Row 7 with Row 4, we can observe that our method achieves more significant performance improvements. The reason is that [11] is designed for time-invariant environments, while our online algorithm can adapt to the time-varying environments in CIL.

Visualizing activation maps. Figure 4 demonstrates the activation maps visualized by Grad-CAM [33] for the final model (obtained after five phases) on ImageNet-Subset 5-phase. The left and right sub-figures show the results for *training-from-half* (TFH) and *training-from-scratch* (TFS), respectively. We can observe: 1) LUCIR [14] makes predictions according to foreground (correct) and background (incorrect) regions in the TFH and TFS settings, respectively; 2) iCaRL [30] behaves opposite to LUCIR in the two settings; 3) our method always makes predictions according to foreground (correct) regions in both TFH and TFS. The reasons are as follows. LUCIR applies a strong (feature) KD by default, so it performs better in TFH. iCaRL applies a weak (logit) KD by default, so it performs better in TFS. Our method can adjust the hyperparameters to change between different KD losses, so it performs well in both settings.

Hyperparameter values. Figure 5 shows the hyperparameter values produced by our policy on CIFAR-100 25-phase.

1) KD loss weights β and γ . From the figure, we have two observations. a) The policy learns to produce a larger initial value for γ and β in TFH and TFS, respectively. Our explanation is as follows. In TFH, we already have a pre-

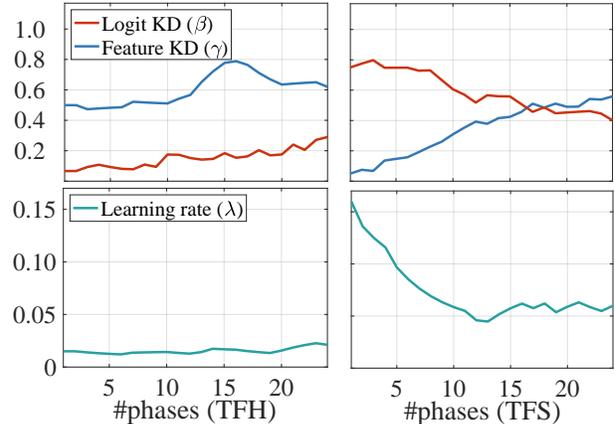


Figure 5. The hyperparameter values produced by our policy on CIFAR-100 25-phase for LUCIR [14] w/ ours. We smooth all curves with a rate of 0.8 for better visualization.

trained model, so we need a strong regularization term (i.e., feature KD loss) to make the model more stable and avoid forgetting. In TFS, we start from random initialization, so we need a weak regularization term (i.e., logit KD loss) to improve the model’s plasticity. b) Both β and γ increase in TFH, while β decreases and γ increases in TFS. It is because we need stronger regularization to maintain the knowledge when more data is observed. The policy achieves that in different ways: it assigns higher weights for both KD losses in TFH, while it transfers from logit KD to feature KD in TFS.

2) Learning rates λ . In Figure 5, we can observe that the learning rate in TFS is much higher than in TFH. It can be explained that 1) we need a higher learning rate in TFS as the model is trained from scratch and needs to capture more new knowledge; 2) we need a lower learning rate in TFH because we need to avoid forgetting the pre-trained model.

3) Classifier type indicator δ . On CIFAR-100 25-phase, our policy learned to choose the NCM and FC classifiers in TFS and TFH, respectively.

6. Conclusions

In this study, we introduce a novel framework that allows us to optimize hyperparameters online to balance the stability and plasticity in CIL. To achieve this, we formulate the CIL task as an online MDP and learn a policy to produce the hyperparameters. Our approach is generic, and it can be easily applied to existing methods to achieve large-margin improvements in both TFS and TFH settings. It is worth mentioning that our method can also be applied to optimize other key hyperparameters in CIL.

Acknowledgments

This research was supported by A*STAR under its AME YIRG Grant (Project No. A20E6c0101).

References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Babak Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *CVPR*, pages 3931–3940, 2020. 2
- [2] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT*, volume 23, pages 39.1–39.26, 2012. 3, 5
- [3] Terry Anderson. *The theory and practice of online learning*. Athabasca University Press, 2008. 2
- [4] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002. 1, 2, 3, 5, 6
- [5] Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010. 3
- [6] Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *JMLR*, 18:1–43, 2018. 3
- [7] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *CVPR*, pages 583–592, 2019. 2
- [8] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, pages 86–102, 2020. 2, 3, 4, 7
- [9] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Experts in a markov decision process. In *NIPS*, pages 401–408, 2005. 3, 5
- [10] Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009. 2, 3, 5, 6
- [11] Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *ICML*, pages 1568–1577, 2018. 3, 7, 8
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [13] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *arXiv*, 1503.02531, 2015. 4
- [14] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [15] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *CVPR*, pages 3957–3966, 2021. 2, 3
- [16] K. J. Joseph, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Vineeth N Balasubramanian. Energy-based latent aligner for incremental learning. In *CVPR*, pages 7452–7461, 2022. 2
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 6
- [18] Yingying Li, Subhro Das, and Na Li. Online optimal control with affine constraints. In *AAAI*, pages 8527–8537, 2021. 3
- [19] Yingying Li and Na Li. Online learning for markov decision processes in nonstationary environments: A dynamic regret analysis. In *2019 American Control Conference (ACC)*, pages 1232–1237. IEEE, 2019. 3
- [20] Yingying Li, Aoxiao Zhong, Guannan Qu, and Na Li. Online markov decision processes with time-varying transition probabilities and rewards. In *ICML workshop on Real-world Sequential Decision Making*, 2019. 3
- [21] Zhizhong Li and Derek Hoiem. Learning without forgetting. In *ECCV*, pages 614–629, 2016. 1, 2, 4, 5
- [22] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, pages 2544–2553, 2021. 1, 2, 4, 5, 6, 7
- [23] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Rmm: Reinforced memory management for class-incremental learning. In *NeurIPS*, pages 3478–3490, 2021. 1, 2, 3, 5, 6, 7, 8
- [24] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, pages 12245–12254, 2020. 2, 3, 5, 6
- [25] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*, volume 24, pages 109–165. Elsevier, 1989. 1
- [26] K. McRae and P. Hetherington. Catastrophic interference is eliminated in pre-trained networks. In *CogSci*, 1993. 1
- [27] Mozghan PourKeshavarzi, Guoying Zhao, and Mohammad Sabokrou. Looking back on learned experiences for class/task incremental learning. In *ICLR*, 2022. 2
- [28] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, pages 524–540, 2020. 2
- [29] R. Ratcliff. Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, 97:285–308, 1990. 1
- [30] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, pages 5533–5542, 2017. 1, 2, 3, 4, 5, 6, 7, 8
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 6
- [32] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv*, 1606.04671, 2016. 2

- [33] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *CVPR*, pages 618–626, 2017. [6](#), [8](#)
- [34] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NIPS*, pages 2990–2999, 2017. [2](#)
- [35] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NIPS*, pages 4077–4087, 2017. [5](#)
- [36] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. In *ECCV*, pages 254–270, 2020. [2](#)
- [37] Amal Rannen Triki, Rahaf Aljundi, Matthew B. Blaschko, and Tinne Tuytelaars. Encoder based lifelong learning. In *ICCV*, pages 1329–1337, 2017. [2](#)
- [38] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, pages 398–414, 2022. [2](#), [7](#)
- [39] Ju Xu and Zhanxing Zhu. Reinforced continual learning. In *NeurIPS*, pages 899–908, 2018. [2](#), [3](#)
- [40] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021. [2](#), [7](#)
- [41] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *CVPR*, pages 6982–6991, 2020. [2](#)